arXiv:1702.04825v2 [cs.LG] 17 Feb 2017

Learning to Use Learners' Advice

Adish Singla ADISH.SINGLA@INF.ETHZ.CH
Hamed Hassani HAMED@INF.ETHZ.CH
Andreas Krause KRAUSEA@ETHZ.CH
ETH Zurich, Zurich, Switzerland

Abstract

In this paper, we study a variant of the framework of online learning using expert advice with limited/bandit feedback. We consider each expert as a learning entity, seeking to more accurately reflecting certain real-world applications. In our setting, the feedback at any time t is limited in a sense that it is only available to the expert i^t that has been selected by the central algorithm (forecaster), *i.e.*, only the expert i^t receives feedback from the environment and gets to learn at time t. We consider a generic black-box approach whereby the forecaster does not control or know the learning dynamics of the experts apart from knowing the following no-regret learning property: the average regret of any expert j vanishes at a rate of at least $\mathcal{O}(t_j^{\beta-1})$ with t_j learning steps where $\beta \in [0,1]$ is a parameter.

In the spirit of competing against the best action in hindsight in multi-armed bandits problem, our goal here is to be competitive w.r.t. the cumulative losses the algorithm could receive by following the policy of always selecting one expert. We prove the following hardness result: without any coordination between the forecaster and the experts, it is impossible to design a forecaster achieving no-regret guarantees. In order to circumvent this hardness result, we consider a practical assumption allowing the forecaster to "guide" the learning process of the experts by filtering/blocking some of the feedbacks observed by them from the environment, *i.e.*, not allowing the selected expert i^t to learn at time t for some time steps. Then, we design a novel no-regret learning algorithm LEARNEXP for this problem setting by carefully guiding the feedbacks observed by experts. We prove that LEARNEXP achieves the worst-case expected cumulative regret of $\mathcal{O}(T^{\frac{1}{2-\beta}})$ after T time steps and matches the regret bound of $\Theta(T^{\frac{1}{2}})$ for the special case of multi-armed bandits.

1. Introduction

Many real-world applications involve repeatedly making decisions under uncertainty—for instance, choosing one of the several items to recommend to the user, dynamically allocating resources among available stock options in a financial market, or sequentially deciding the next medical test in healthcare. Furthermore, the feedback is often limited in these settings in a sense that only the loss/reward associated with the action taken by the system is observed, referred to as the bandit feedback setting. Online learning using expert advice with bandit/limited feedback is a well-studied framework to model the above-mentioned application settings (Freund and Schapire, 1995; Auer et al., 2002; Cesa-Bianchi and Lugosi, 2006; Bubeck and Cesa-Bianchi, 2012) and addresses the fundamental question of how a learning algorithm should trade-off exploration (the cost of acquiring new information) versus exploitation (acting greedily based on current information to minimize instantaneous losses). In this paper, we investigate this framework with an important practical consideration:

How do we use the advice of experts when they themselves are learning entities?

© A. Singla, H. Hassani & A. Krause.

1.1. Motivating Applications

Modeling experts as learning entities realistically captures many practical scenarios of how one would define/encounter these experts in real-world applications, such as seeking advice from fellow players or friends, aggregating prediction recommendations from trading agents or different marketplaces, product testing with human participants who might adapt over time, information acquisition from crowdsourcing participants who might learn over time, the problem of meta-learning and hyperparameter tuning whereby different learning algorithms are treated as experts (*cf.* Baram et al. (2004); Hsu and Lin (2015)), and many more.

As a concrete running example, we consider the problem of learning to offer personalized deals / discount coupons to users enabling new businesses to incentivize and attract more customers (Edelman et al., 2011; Singla et al., 2016). An emerging trend is *deal-aggregator* sites like *Yipit*[1] providing personalized coupon recommendation services to their users by aggregating and selecting coupons from *daily-deal* marketplaces like *Groupon* and *LivingSocial*[1]. One of the primary goals of these recommendation systems like *Yipit* (corresponding to the central algorithm / forecaster in our setting) is to design better selection strategies for choosing coupons from different marketplaces (corresponding to the experts in our setting). However, these marketplaces (experts) themselves would be learning to optimize the coupons to offer, for instance, the discount price or the type of the coupon based on historic interactions with users (Edelman et al., 2011).

1.2. Experts as Learning Entities: Challenges and Our Results

We now provide an overview of our approach, the main challenges in designing a forecaster with no-regret guarantees, and our results.

The interaction model. We consider an online setting similar to that of adversarial online learning using experts' advice with bandit feedback (Auer et al., 2002). However, to keep the presentation more general (*e.g.*, we do not necessarily require that the sets of actions are shared across different experts), the forecaster selects / seeks advice from only one expert i^t at any time t (*cf.* Kale (2014) for more discussion about this aspect). More specifically, at time t, the forecaster selects an expert i^t, performs an action $a^t_{i^t}$ recommended by the expert i^t, and incurs a loss $l^t(a^t_{i^t})$ set by the adversary.

The notion of regret. In the standard framework, *i.e.*, when the experts are not learning entities, the EXP3 algorithm (Auer et al., 2002) is well-suited for this problem setting achieving the optimal regret bounds. However, it is important to note that the classical notion of *external* regret used in the literature (*cf.* Auer et al. (2002); Cesa-Bianchi and Lugosi (2006); Bubeck and Cesa-Bianchi (2012)) does not provide any meaningful guarantees in our setting in terms of competing against the "best" expert. Similar to the notion of competing against the best action in hindsight in multi-armed bandits problem, we want to be competitive w.r.t. the cumulative losses the algorithm could receive by following the policy of always selecting one expert (*cf.* Section 2.3 for a formal definition).

Experts as no-regret learners and blackbox approach. In our setting, the experts themselves are learning entities. Formally, we assume that the experts are no-regret learners, *i.e.*, the average regret of any expert j vanishes at a rate of *at least* $\mathcal{O}(t_j^{\beta-1})$ with t_j learning steps where $\beta \in [0, 1]$ is a parameter known to the forecaster. We consider the following natural notion of bandit/limited feedback: only the selected expert i^t receives feedback and gets to learn at time t; all other experts that have not been selected at time t experience no change in their learning state at this time. We

1. http://yipit.com/; http://www.groupon.com/; https://livingsocial.com/

consider a generic black-box approach in which the forecaster does not know and cannot control the internal learning dynamics of the experts.

Challenges and hardness result. It turns out that modeling these experts as learning entities leads to a challenging twist in this well-studied and foundational online learning framework. In this paper, we prove the following hardness result for our problem setting: without any coordination between the forecaster and the experts, it is impossible to design a forecaster achieving no-regret guarantees in the worst-case. Somewhat surprisingly, this hardness result holds when playing against an oblivious (non-adaptive) adversary and even if restricting the experts to be implementing some well-studied online learning algorithms, for instance, the HEDGE algorithm (Freund and Schapire, 1995). The fundamental challenge leading to this hardness result arises from the fact that the forecaster's selection strategy affects the feedback sequences observed by the experts which in turn alters their learning process.

"Guided" feedbacks and achieving no-regret guarantees. In order to circumvent this hardness result, we consider the following practical assumption: we allow the forecaster to "guide" the learning process of the experts by filtering/blocking some of the feedbacks observed by them from the environment, *i.e.*, the selected expert i^t would not learn at time t for some time steps. For instance, in the motivating application of offering personalized deals to users, the deal-aggregator site (forecaster) often primarily interacts with users on behalf of the individual daily-deal marketplaces (experts) and hence can control the flow of feedback to these marketplaces. Alternatively, we note that this process of guiding and restricting the feedback can be achieved via coordination between the forecaster and the selected expert i^t with a 1-bit of communication at time t. Given this additional control, we design a novel algorithm LEARNEXP for the forecaster which carefully guides the feedbacks observed by experts. We prove that LEARNEXP achieves the worst-case expected cumulative regret of $\mathcal{O}(T^{\frac{1}{2-\beta}})$ after T time steps against an oblivious adversary for a rich family of no-regret learning algorithms that experts may be implementing. For the special case of multi-armed bandits, algorithm LEARNEXP is equivalent to that of the well-studied EXP3 algorithm and hence matches the optimal regret bound of $\Theta(T^{\frac{1}{2}})$.

Connections to the existing results. Maillard and Munos (2011) studied the problem of competing against an adaptive adversary when the adversary's reward generation policy is restricted to a pre-specified set of known models. For this problem, the authors introduced the EXP4/EXP3 algorithm, *i.e.*, EXP4 meta-algorithm with experts executing EXP3 algorithms proving a regret of $\mathcal{O}(T^{\frac{2}{3}})$ (*cf.* Bubeck and Cesa-Bianchi (2012) for a variant of the algorithm). This EXP4/EXP3 algorithm is perhaps closest to ours, as it involves a forecaster where the experts are the learning entities. However, we note that our hardness result does not contradict their regret bounds—the key difference in their setting is that the forecaster has the power to modify the losses as seen by experts, and it provides an unbiased estimate of the losses to these experts. Moreover, their analysis is specific to the experts implementing the EXP3 or bandit algorithms, whereas the focus of this paper is to present a more generic learning framework in which experts as learning entities may implement a broad class of learning algorithms. Our work is also related to contemporary work by Agarwal et al. (2016), who study a variant of the problem tackled by Maillard and Munos (2011) and also prove a hardness result similar to that of ours. Note that, in applications where the experts directly receive feedback from the environment, implementing the strategies of Maillard and Munos (2011); Agarwal et al. (2016) would require the forecaster to communicate the probability p with which the expert i^t was selected at time t. However, our proposed idea of guiding the feedback

Protocol 1: The interaction between adversary ADV, algorithm ALGO, and experts

foreach $t = 1, 2, \ldots, T$ **do**

/* Adversary generates the following */

a private loss vector l^t, *i.e.*, $l^t(a)\ \forall\ a \in \mathcal{A}$

a private feedback vector f^t, *i.e.*, $f^t(a)\ \forall\ a \in \mathcal{A}$

a public context $x^t \in \mathcal{X}$

/* Selecting an expert and performing an action */

ALGO selects an expert $i^t \in [N]$ denoted as EXP_{i^t}

ALGO performs the action $a^t_{i^t}$ recommended by EXP_{i^t}

/* Feedback and updates */

ALGO incurs (and observes) loss $l^t(a^t_{i^t})$ and updates its selection strategy

$\forall j \in [N] : j \neq i^t$, EXP_j does not observe any feedback and makes no update

EXP_{i^t} observes feedback $f^t(a^t_{i^t})$ from the environment and updates its learning state

end

can be achieved via coordination between the forecaster and the selected expert i^t with a 1-bit of communication at time t.

2. The Model

We have the following entities in our problem setting: (i) an algorithm ALGO as the forecaster; (ii) the adversary ADV acting on behalf of the environment; and (iii) N experts $\text{EXP}_j\ \forall j \in \{1, \ldots N\}$ (henceforth denoted as $[N]$).

2.1. Specification of the Interaction

Protocol 1 provides a high-level specification of the interaction between the $N + 2$ entities. The sequential decision making process proceeds in rounds $t = 1, 2, \ldots, T$ (henceforth denoted as $[T]$); for simplicity we assume that T is known in advance to the algorithm and the results in this paper can be extended to an unknown horizon via the usual doubling trick (Cesa-Bianchi and Lugosi, 2006). Each expert EXP_j where $j \in N$ is associated with a set of actions $\mathcal{A}_j$ and the action set of the algorithm ALGO is given by $\mathcal{A} = \cup_{j \in [N]} \mathcal{A}_j$. For the clarify of presentation in defining the loss and feedback vectors, we will consider that the action sets of experts are disjoint.[2]

At any time t, the adversary ADV generates a private loss vector l^t (*i.e.*, $l^t(a)\ \forall\ a \in \mathcal{A}$) and a private feedback vector f^t (*i.e.*, $f^t(a)\ \forall\ a \in \mathcal{A}$). Additionally, the adversary ADV generates a context $x^t \in \mathcal{X}$ that is accessible to all the experts while recommending their actions at time t—this context essentially encodes all the side information from the environment accessible to the experts at time t (*e.g.*, this context could represent preferences of a user arriving at time t in an online recommendation system). Simultaneously, the algorithm ALGO (possibly with some randomization) selects expert EXP_{i^t} to seek advice. The selected expert EXP_{i^t} recommends an action $a_{i^t} \in \mathcal{A}_{i^t} \subseteq \mathcal{A}$ (possibly with its internal randomization) which is then performed by the algorithm. As feedback, the algorithm ALGO observes the loss $l^t(a^t_{i^t})$ and updates its strategy on how to select experts in the future. All the experts apart from the one selected (*i.e.*, $\text{EXP}_j\ \forall\ j \neq i^t$) observe no feedback and

2. Note that assuming the disjoint action sets across experts is w.l.o.g., as we can still simulate the shared actions by enforcing a constraint that the losses generated by the adversary are same for the shared actions at any given time.

make no update at this time. The selected expert EXP_{i^t} observes a feedback from the environment denoted as $f^t(a^t_{i^t})$ and updates its learning state. At the end of time t, the algorithm ALGO incurs a loss of $l^t(a^t_{i^t})$.

So far, we have considered a generic notion of the feedback received by the selected expert—this feedback essentially depends on the application setting and is supposed to be "compatible" with the learning algorithm used by an expert. As a concrete example, consider an expert EXP_j implementing the EXP3 algorithm and taking action a^t_j at time t, then the feedback $f^t(a^t_j)$ received by this expert (if selected at time t) is the loss $l^t(a^t_j)$; for the case of expert EXP_j implementing the HEDGE algorithm, the feedback $f^t(a^t_j)$ received by this expert (if selected at time t) is the set of losses $l^t(a)\ \forall\ a \in \mathcal{A}_j$. The feedback could be more general, for instance, receiving a binary signal of rejection or acceptance of the offered deal when an expert is implementing a dynamic pricing based algorithm via the partial monitoring framework (Cesa-Bianchi and Lugosi, 2006; Bartók et al., 2014). Also, we note that the special case of standard multi-armed bandits is captured by the setting in which $\mathcal{A}_j$ is a singleton for every expert $j \in [N]$.

We assume that the losses are bounded in the range $[0, l_{max}]$ for some known $l_{max} \in \mathbb{R}_+$; w.l.o.g. we will use $l_{max} = 1$ (Auer et al., 2002). We consider an oblivious (non-adaptive adversary) as is usual in the literature (Freund and Schapire, 1995; Auer et al., 2002), *i.e.*, the loss vector l^t, the feedback vector f^t, and the context x^t at any time t do not depend on the actions taken by ALGO, and hence can be considered to be fixed in advance. Apart from that, no other restrictions are put on the adversary, and it has complete knowledge about the algorithm ALGO and the learning dynamics of the experts.

2.2. Specification of the Experts

We consider a generic black-box approach in which ALGO does not know and cannot control the internal dynamics of the experts. In order to formally state the objective and guarantees we seek, we now provide a generic specification of the experts. At time t, let us denote an instance of feedback received by EXP_{i^t} by a tuple $h = (a^t_{i^t}, x^t, f^t(a^t_{i^t}))$. For any expert EXP_j where $j \in [N]$, let $\mathcal{H}^t_j = (h^1, h^2, \ldots)$ denote the feedback history for EXP_j, *i.e.*, an ordered sequence of feedback instances observed by EXP_j up to time t. The length $|\mathcal{H}^t_j|$ denotes the number of learning steps for EXP_j up to time t. At time t, the action a^t_j recommended by EXP_j to the algorithm, if this expert is selected, is given by $a^t_j = \pi_j(x^t, \mathcal{H}^t_j)$ where π_j is a (possibly randomized) function of EXP_j, taking as input a context and a history of feedback sequence, and outputs an action $a \in \mathcal{A}_j$. Importantly, this history $\mathcal{H}^t_j$ is dependent on the execution of the algorithm ALGO— for clarify of presentation, we denote it as $\mathcal{H}^t_{j,\text{ALGO}}$.

No-regret learning dynamics. To be able to say anything meaningful in this setting, we introduce the constraint of *no-regret* learning dynamics on the experts.[3] Let us consider any sequence of loss vector l, feedback vector f, and context x given by $\mathcal{S} = \big((l^\tau, f^\tau, x^\tau)\big)_{\tau=\{1,2,\ldots\}}$ generated arbitrarily by ADV and let $|\mathcal{S}|$ denotes its length. Consider a setting in which an expert EXP_j for any $j \in [N]$ is selected at every time step. At every time step $\tau \in [|\mathcal{S}|]$, EXP_j recommends an action a^τ_j, accumulates the loss $l(a^\tau_j)$, and observes the feedback $f(a^\tau_j)$. In this setting, EXP_j observes feedback at every time step and we denote this "complete" history of feedback sequence at any time $\tau \in |\mathcal{S}|$ as $\mathcal{H}^t_{j,1}$ whereby 1 denotes the fact that this expert is selected and receives feedback with

3. In order to prove the no-regret guarantees for our algorithm LEARNEXP in Section 4, this constraint is required to hold *only* for the *best* expert against which we want to be competitive, a less stringent requirement.

probability 1 at every time step. Then, the no-regret learning dynamics of EXP_j parameterized by $\beta_j \in [0,1]$ guarantees that the expected average regret vanishes as follows[4]:

$$\mathbb{E}\left[\frac{1}{|\mathcal{S}|}\sum_{\tau=1}^{|\mathcal{S}|} l^\tau\left(\pi_j(x^\tau, \mathcal{H}^\tau_{j,1})\right)\right] - \mathbb{E}\left[\frac{1}{|\mathcal{S}|}\sum_{\tau=1}^{|\mathcal{S}|} l^\tau\left(\pi_j(x^\tau, \mathcal{H}^{|\mathcal{S}|}_{j,1})\right)\right] \leq \mathcal{O}(|\mathcal{S}|^{\beta_j - 1}) \tag{1}$$

where the expectation is w.r.t. the randomization of function π_j. We assume that parameter $\beta \in [0,1]$ upper bounds the regret rate parameters of individual experts and is a parameter known to the forecaster.[5]

2.3. Our Objective: No-Regret Guarantees

Intuitively, we want to be competitive w.r.t. the cumulative losses the algorithm could receive by following the policy of always using the advice of one single expert—such a policy ensures that the single expert gets more feedback to improve its learning state and hence incur less cumulative loss. This is a challenging problem when the experts are learning entities. For instance, what may go wrong is that the *best* expert could have a slow rate of learning/convergence thus incurring high losses in the beginning, misleading the algorithm to essentially "downweigh" this expert. This is turn further exacerbates the problem for the best expert in the bandit feedback setting as this expert will be selected less and will have fewer learning steps to improve its state. This adds new challenges to the classic trade-off between exploration and exploitation, suggesting the need to explore at higher rate to tackle this problem.

Let us begin by looking at the classic notion of *external* regret used in the literature (Auer et al., 2002; Cesa-Bianchi and Lugosi, 2006; Bubeck and Cesa-Bianchi, 2012). Given that the experts are learning entities, naturally the losses incurred at any time step are dependent on the history of the forecaster's actions as that history defines the current learning state of the individual experts. Given this subtle issue of history dependent losses, the usual notion of *external* regret does not provide any meaningful guarantees in terms of competing against the "best expert in hindsight" (see below for a formal definition); the bounds given by the *external* regret are only w.r.t. the post hoc sequence of actions performed and losses observed during the execution of the algorithm (*cf.* Maillard and Munos (2011); Arora et al. (2012); McMahan and Streeter (2009) for more discussion on this).

We consider the following natural notion of regret in this paper: our goal is to be competitive w.r.t. the *best expert in hindsight*, that is, competitive w.r.t. the cumulative loss that any expert could have received with the optimal actions it could have taken in hindsight. Formally, the expected cumulative regret of ALGO against the best expert in hindsight is given by:

$$\text{REG}(T, \text{ALGO}) := \sum_{t=1}^{T} \mathbb{E}\left[l^t\left(\pi_{i^t}(x^t, \mathcal{H}^t_{i^t,\text{ALGO}})\right)\right] - \min_{j\in[N]} \mathbb{E}\left[\sum_{t=1}^{T} l^t\left(\pi_j(x^t, \mathcal{H}^T_{j,1})\right)\right] \tag{2}$$

where the expectation is w.r.t. the randomization of the algorithm as well as any internal randomization of the experts. Our goal is to design an algorithm ALGO for the forecaster so that the regret $\text{REG}(T, \text{ALGO})$ grows sublinearly in time T.

4. Note that this is a weaker notion of regret—any deterministic policy π_j that always outputs a constant action has $\beta_j = 0$. However, $\beta_j = 0$ would be the right way to characterize the learning dynamics of this expert for our setting.
5. Again for our algorithm LEARNEXP in Section 4, β *only* needs to upper bound the regret rate for the *best* expert against which we want to be competitive.

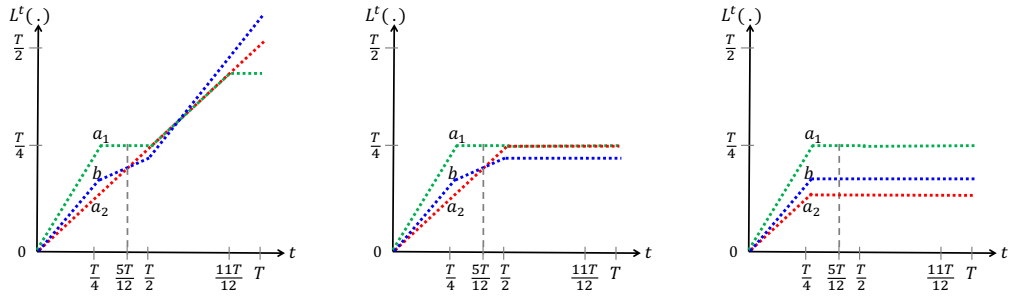

(a) Cumulative losses (L_1) (b) Cumulative losses (L_2) (c) Cumulative losses (L_3)

Figure 1: We have two experts: EXP$_1$ plays HEDGE and has two actions $\mathcal{A}_1 = \{a_1, a_2\}$, and EXP$_2$ has only one action $\mathcal{A}_2 = \{b\}$. Figures 1(a), 1(b), and 1(c) shows the *cumulative* loss sequences L_1, L_2, and L_3 for three different scenarios—the adversary at $t = 0$ uniformly at random picks one of these scenarios and uses that loss sequence. These plots show the cumulative losses of the three actions $\mathcal{A} = \{a_1, a_2, b\}$ for three different sequences. The losses are illustrated with the following color scheme—a_1:*green*, a_2:*red*, and b:*blue*.

3. Hardness Result

We show in this section that, in the absence of any coordination between the forecaster and experts, it is impossible to design a forecaster that achieves no-regret guarantees in the worst-case. Somewhat surprisingly, we prove this hardness result when playing against an oblivious (non-adaptive) adversary and when restricting the experts to be implementing the well-studied HEDGE algorithm (Freund and Schapire, 1995). We formally state this hardness result in the Theorem 1 below.

Theorem 1 *There is a setting in which each of the experts has no-regret learning dynamics with parameter $\beta = \frac{1}{2}$; however, any algorithm* ALGO *(forecaster) will suffer a positive average regret, i.e.,* $\text{REG}(T, \text{ALGO}) = \Omega(T)$.

The proof is given in the Appendix, we briefly outline the main ideas below. Our setting for proving this theorem consists of two experts EXP$_1$ and EXP$_2$. The first expert EXP$_1$ has two actions given by $\mathcal{A}_1 = \{a_1, a_2\}$, and the second expert EXP$_1$ has only one action given by $\mathcal{A}_2 = \{b\}$. The action set of the algorithm ALGO is given by $\mathcal{A} = \{a_1, a_2, b\}$. The expert EXP$_1$ plays the HEDGE algorithm (Freund and Schapire, 1995), *i.e.*, the regret rate parameter is $\beta_1 = 0.5$; the expert EXP$_2$ has only one action to play as in the standard multi-armed bandit with $\beta_2 = 0$.[6] Figures 1(a), 1(b), and 1(c) show the *cumulative* loss sequences L_1, L_2, and L_3 for three different scenarios—the adversary at $t = 0$ uniformly at random picks one of these scenarios and uses that loss sequence.

The main idea of the proof uses the following arguments. We consider the case where the forecaster is facing the sequence L_1 (chosen by the adversary with probability $\frac{1}{3}$ at $t = 0$). We then divide the time horizon T into different slots and discuss the execution behavior of the forecaster and experts over these time slots. Specifically, our claim is that in the time slot $t \in (\frac{T}{4}, \frac{T}{2}]$, the expert

6. In fact, this hardness result holds even when considering a powerful forecaster which knows exactly the learning algorithms used by the experts, and is able to see the losses $\{l^t(a_1), l^t(a_2), l^t(a_3)\}$ at every time $t \in [T]$.

EXP_1 would not be selected for $\frac{T}{12} - o(T))$ time steps. As a result, in the time slot $t \in (\frac{11T}{12}, T]$, the expert EXP_1 would select action a_2 almost surely, and a_1 would only be selected $o(T)$ number of times, leading to a positive average regret for the forecaster. Informally speaking, our negative example shows that the forecaster's selection strategy could add "blind spots" in the feedback history seen by the experts and that they might not be able to "recover" from this. The key fundamental challenge leading to this hardness result is that the forecaster's selection strategy affects the feedback sequences observed by the experts, which in turn alters the experts' learning process.

4. Our Algorithm LEARNEXP

In this section, we introduce a practical assumption that allows the forecaster to "guide" the learning process of experts, and then we design our main algorithm LEARNEXP with provable no-regret guarantees.

4.1. Guided Feedbacks

In order to circumvent the hardness result proved in Section 3, we now consider a practical assumption motivated by the application setting of deal-aggregator sites, as discussed in Section 1. Usually, a deal-aggregator site interacts with users on behalf of the individual daily-deal marketplaces (experts) and hence could control the flow of feedback to these marketplaces. Hence, we allow the forecaster to "guide" the learning process of the experts by filtering/blocking some of the feedback the experts receive from the environment. Recall that at time t, as per the interaction model presented in Section 2, the selected expert EXP_{i^t} observes feedback $f^t(a^t_{i^t})$ from the environment. We now consider the setting with the following additional power in the hands of the forecaster: In order to guide the learning process of the experts, the forecaster at time t could block the feedback, *i.e.*, the expert EXP_{i^t} would not observe feedback at time t and hence would not learn at this time (just like any other experts who were not selected at time t). Alternatively, we note that this process of guiding the feedback could be achieved via coordination between the forecaster and the selected expert EXP_{i^t} with a 1-bit communication at time t.

4.2. Algorithm LEARNEXP

With this additional power of the forecaster to guide feedback, we develop our main algorithm LEARNEXP, presented in Algorithm 2. The selection strategy of the algorithm LEARNEXP is similar to the EXP family of algorithms, and in particular is equivalent to the EXP3 algorithm by Auer et al. (2002). The core idea of guiding the feedbacks observed by experts is presented in Lines 10,11, and 12.

By default, as per the Protocol 1, the selected expert EXP_{i^t} always observes feedback at time t—for this protocol, the hardness result of Theorem 1 applies. Our algorithm LEARNEXP instead decides whether the expert EXP_{i^t} should observe/use the feedback based on the outcome ξ^t of a coin flip with probability $\frac{\eta}{N \cdot p^t_{i^t}}$. By choosing this particular probability, the algorithm LEARNEXP ensures that the probability that any expert EXP_j observes feedback at time t is constant over time and is given by $\frac{\eta}{N}$. The key parameter of the algorithm η would be fixed in Theorem 3 based on the regret rate β to achieve the desired guarantees on the regret.

The guarantees in Theorem 3 mean that by adding this additional control/coordination in our model, we are able to circumvent the hardness result of Theorem 1. Interestingly, if we consider any

Algorithm 2: LEARNEXP

Parameters: $\eta \in (0, 1]$
Initialize: time $t = 1$, weights $w_j^t = 1 \; \forall j \in [N]$
foreach $t = 1, 2, \ldots, T$ **do**
  /* Selecting an expert and performing an action */
  $\forall j \in [N]$, define probability $p_j^t = (1 - \eta) \cdot \frac{w_j^t}{\left(\sum_{k \in [N]} w_k^t\right)} + \frac{\eta}{N}$
  Draw i^t from the multinomial distribution $(p_j^t)_{j \in [N]}$
  Perform action $a_{i^t}^t$ recommended by the expert EXP_{i^t}
  /* Observing the loss and making updates */
  Observe loss $l^t(a_{i^t}^t)$
  $\forall j \in [N]$, do the following:
    Set $\widetilde{l}_j^t$ as follows: $\widetilde{l}_j^t = \frac{l^t(a_{i^t}^t)}{p_{i^t}^t}$ for $j = i^t$, else $\widetilde{l}_j^t = 0$
    Update $w_j^{t+1} \leftarrow w_j^t \cdot \exp(-\frac{\eta \cdot \widetilde{l}_j^t}{N})$
  /* Guiding the feedback */
  $\xi^t \sim Bernoulli(\frac{\eta}{N \cdot p_{i^t}^t})$
  if $(\xi^t = 1)$ **then**
    EXP_{i^t} observes feedback $f^t(a_{i^t}^t)$ from the environment and updates its learning state
  end
end

expert EXP_j for $j \in [N]$, the history $\mathcal{H}_j^t$ at any time t under this guided feedback setting would only contain a subset of the feedback instances that it would have received without guiding (*i.e.*, where $\xi^t = 1 \; \forall t \in [T]$). By carefully allowing the expert to observe a strictly smaller set of feedback instances allows us to ensure that the expert EXP_j achieves low regret. Considering the example we use in the proof of Theorem 1 to show the hardness results, this means that by carefully guiding the feedback received by experts, our algorithm LEARNEXP ensures that there are no "blind spots" in the feedback history of any expert. However, in order to achieve this, the algorithm is required to explore at a higher rate, as is evident by the value of η in Theorem 3.

4.3. Theoretical Guarantees

Next, we analyze the theoretical guarantees of our algorithm LEARNEXP. One approach to doing this is to consider a particular class of no-regret learning algorithms that experts implement and prove guarantees for that class. Instead, we introduce a novel, generic notion of "smooth" no-regret learning—our theoretical guarantees are then proven for the experts that have *no-regret* and *smooth* learning dynamics. Next, we introduce this notion and then discuss (see Proposition 2) the class of no-regret learning algorithms that also satisfy the constraint of smooth learning dynamics.

4.3.1. SMOOTH LEARNING DYNAMICS

In our bandit feedback setting, not all the experts can observe feedback at a given time step, and hence the history of feedback instances received by any particular expert is naturally "sparse". To formally state the behavior of the learning algorithm under this sparse feedback, we now introduce a new notion, termed *smooth* learning dynamics, to complement the no-regret learning dynamics defined in (1). Consider the same fixed sequence $\mathcal{S}$ as used in defining (1) and an expert EXP_j. However, instead of observing feedback at every time step, let's say that the expert EXP_j only gets to observe the feedback sporadically at a rate of $\alpha \in (0, 1]$—we call this an α-sparse history, denoted as $\mathcal{H}^l_{j,\alpha}$. Then, the constraint of smooth learning dynamics ensures that the expected regret of the expert EXP_j when receiving the above-mentioned sparse feedback vanishes (smoothly w.r.t. rate α) as follows:

$$\mathbb{E}\left[\frac{1}{|\mathcal{S}|}\sum_{\tau=1}^{|\mathcal{S}|} l^\tau\left(\pi_j(x^\tau, \mathcal{H}^\tau_{j,\alpha})\right)\right] - \mathbb{E}\left[\frac{1}{|\mathcal{S}|}\sum_{\tau=1}^{|\mathcal{S}|} l^\tau\left(\pi_j(x^\tau, \mathcal{H}^{|\mathcal{S}|}_{j,1})\right)\right] \leq \mathcal{O}((\alpha \cdot |\mathcal{S}|)^{\beta_j - 1}) \tag{3}$$

where the expectation is w.r.t. the randomization of function π_j as well as w.r.t. the randomization in generating this sparse history. The following proposition states that a rich class of online learning algorithms indeed have *smooth* learning dynamics that can be used by the experts, *cf.* Appendix for the proof.

Proposition 2 *A rich class of no-regret online learning algorithms based on gradient-descent style updates have smooth learning dynamics including the Online Mirror Descent family of algorithms with exact or estimated gradients (Shalev-Shwartz, 2011) and Online Convex Programming via greedy projections (Zinkevich, 2003).*

4.3.2. NO-REGRET GUARANTEES OF LEARNEXP

Next, we prove the no-regret guarantees of our algorithm LEARNEXP, formally stated in Theorem 3. The following theorem (stating only the leading terms w.r.t. the T and dropping any other constants like N) provides the no-regret guarantees of LEARNEXP against the best expert in hindsight as per (2). The proof is given in the Appendix.

Theorem 3 *Let T be the fixed time horizon. Consider that the best expert $j^* \in [N]$ has no-regret smooth learning dynamics parameterized by $\beta_{j^*} \in [0, 1]$ and* LEARNEXP *is invoked with input $\beta \in [0, 1]$ such that $\beta \geq \beta_{j^*}$. Set parameters $\eta = \Theta\big(T^{-\frac{1-\beta}{2-\beta}} \cdot N^{\frac{1-\beta}{2-\beta}} \cdot (\log N)^{(\frac{1}{2}\cdot \mathbf{1}_{\{\beta=0\}})}\big)$. Then, for sufficiently large T, the worst-case expected cumulative regret of* LEARNEXP *against the best expert in hindsight is:*

$$\text{REG}(T, \text{LEARNEXP}) \leq \mathcal{O}\big(T^{\frac{1}{2-\beta}} \cdot N^{\frac{1}{2-\beta}} \cdot (\log N)^{(\frac{1}{2}\cdot \mathbf{1}_{\{\beta=0\}})}\big)$$

For the special case of multi-armed bandits (where $\beta = 0$), this regret bound matches the bound of $\Theta(T^{\frac{1}{2}})$—in fact, for this special case, our algorithm LEARNEXP is exactly equivalent to EXP3. For an important case when experts are implementing algorithms like HEDGE or EXP3 (where $\beta = \frac{1}{2}$), our algorithm LEARNEXP achieves the bound of $\mathcal{O}(T^{\frac{2}{3}})$.

5. Background and Related Work

In this section, we provide an overview of the relevant literature.

5.1. Background

We begin with a background on the framework of learning using expert advice with bandit feedback.

Using expert advice. The seminal work of Littlestone and Warmuth (1994); Cesa-Bianchi et al. (1997) initiated the study of using expert advice for prediction problems, and Freund and Schapire (1995) introduced the algorithm HEDGE for the general problem of dynamically allocating resources among a set of options using expert advice.

Using expert advice with bandit feedback. However, the feedback is often limited in these settings in a sense that only the loss/reward associated with the action taken by the system is observed, referred to as the bandit feedback setting. To tackle this, Auer et al. (2002) extended this framework to the limited feedback setting and introduced the EXP family of algorithms (EXP3, EXP4, and its variants) for multi-armed bandits and expert advice with bandit feedback. With limited feedback, this framework addresses the fundamental question of how a learning algorithm should trade-off *exploration* versus *exploitation*. This framework has been studied extensively by researchers in a variety of fields and the above mentioned algorithms provide minimax optimal no-regret guarantees—we refer the reader to Bubeck and Cesa-Bianchi (2012) for the survey on bandit problems and monograph by Cesa-Bianchi and Lugosi (2006).

Furthermore, this framework is very generic and versatile to capture many complex real-world scenarios. For instance, in the EXP family of algorithms (Auer et al., 2002; McMahan and Streeter, 2009; Beygelzimer et al., 2011), each expert could have access to an arbitrary context (*e.g.*, information about user preferences) that may not be shared among experts and may not be available to the algorithm. Furthermore, no statistical assumptions are needed on the process generating context or losses/rewards over time. Consequently, this framework has been used in many diverse application settings including search engine ad placement (McMahan and Streeter, 2009), personalized news article recommendation (Beygelzimer et al., 2011), packet routing in networks, (Awerbuch and Kleinberg, 2004), and meta-learning with different learning algorithms as the experts (Baram et al., 2004; Hsu and Lin, 2015).

5.2. Related Work

Next, we review research work that is relevant to the problem studied in this paper.

Markovian, rested, and restless bandits. The seminal work of Gittins (1979) considered Markovian bandits where each action/arm is associated with its own stochastic MDP and introduces the *Gittins index* to find an optimal sequential policy. Note that the arm changes its state only when it is pulled, hence also termed as *rested* bandits. Whittle (1988) considered an extension termed *restless* bandits where all the arms change their reward distributions at every time step according to their associated stochastic MDP. Restless bandits are notoriously difficult to tackle (*cf.* (Slivkins and Upfal, 2008)), thereby Slivkins and Upfal (2008) considered a type of restless bandits where the change in state is governed by a more gradual process with stochastic rewards depending upon the state. Besbes et al. (2014) considered another type of restless bandits with stochastic reward functions, however these distributions change adversarially with a budget on the allowed variation. Our approach is similar in spirit to the rested bandits; however, none of the frameworks above would model the learning dynamics of the experts in the adversarial setting we consider.

Non-oblivious/adaptive adversary. As in our setting, the challenge of history-dependent expert rewards also arises in the case of non-oblivious/adaptive adversary (Maillard and Munos, 2011; Arora et al., 2012). Arora et al. (2012) studied online learning with bandit feedback against a non-oblivious adversary with *bounded* memory and introduced the notion of policy regret instead of the usual notion of external regret. Maillard and Munos (2011) studied competing against adaptive adversary when the adversary's reward generation policy is restricted to a pre-specified set of known models. However, none of the frameworks of non-oblivious/adaptive adversary listed above model learning dynamics in our setting: It would require an adversary with unbounded memory to apply the results of Arora et al. (2012), and an adversary with unbounded number of models to apply the techniques of Maillard and Munos (2011).

Contextual bandits. Another perspective on tackling some of the applications we mentioned above is the contextual bandit framework (Li et al., 2010; Langford and Zhang, 2007; Agarwal et al., 2014). We refer the reader to the paper by McMahan and Streeter (2009) for more discussion on the connection between the framework of contextual bandits and learning using expert advice with bandit feedback.

Learning in games. An orthogonal line of research studies the interaction of agents in multiplayer games where each agent uses a no-regret learning algorithm (Blum and Monsour, 2007; Syrgkanis et al., 2015). The questions tackled in this line of research are very different as it focuses on the interactions of the agents, their individual as well as social utilities, and the convergence of the game to equilibrium. This orthogonal line of research reassures that the no-regret learning dynamics that we consider in this paper are indeed important and natural dynamics that are also prevalent in other application domains.

6. Conclusions

In this paper, we investigated the online learning framework using expert advice with bandit feedback with an important practical consideration: how do we use the advice of the experts when they themselves are learning entities? As our first contribution, we proved the hardness result stating that it is impossible to achieve no-regret guarantees when the experts receive feedback directly from the environment and there is no further coordination between forecaster/experts. Our hardness result sheds light on the complexity of the problem when applying this online learning framework to real-world applications whereby it is natural for experts to exhibit learning dynamics.

Then, we considered a practical assumption of "guided" feedbacks whereby the forecaster can block/filter the feedback received by the selected expect from the environment. Under this setting, we proposed a novel algorithm LEARNEXP—we proved that LEARNEXP achieves the worst-case expected cumulative regret of $\mathcal{O}(T^{\frac{1}{2-\beta}})$ after T time steps where β is a parameter characterizing the individual no-regret learning dynamics of the best expert. This regret bound matches the bound of $\Theta(T^{\frac{1}{2}})$ for the special case of multi-armed bandits.

There are a number of research directions for future work. An interesting question to tackle is whether it is possible to design a forecaster in our setting with a worst-case cumulative regret of $\Theta(T^{\frac{1}{2}})$ when the individual experts have no-regret learning dynamics with $\beta = \frac{1}{2}$. In this paper, in order to circumvent the hardness result, we considered the power of blocking/filtering the feedbacks, which can equivalently be achieved with a 1-bit of communication at every time step. An interesting direction would be to consider other practical ways of coordination and to understand the minimal coordination required to achieve no-regret guarantees.

References

Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *ICML*, 2014.

Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E. Schapire. Corralling a band of bandit algorithms. *CoRR*, abs/1612.06246, 2016.

Raman Arora, Ofer Dekel, and Ambuj Tewari. Online bandit learning against an adaptive adversary: from regret to policy regret. In *ICML*, 2012.

Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multi-armed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.

Baruch Awerbuch and Robert D Kleinberg. Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches. In *STOC*, 2004.

Yoram Baram, Ran El-Yaniv, and Kobi Luz. Online choice of active learning algorithms. *Journal of Machine Learning Research*, 2004.

Gábor Bartók, Dean P Foster, Dávid Pál, Alexander Rakhlin, and Csaba Szepesvári. Partial monitoring – Classification, regret bounds, and algorithms. *Mathematics of Operations Research*, 2014.

Omar Besbes, Yonatan Gur, and Assaf J. Zeevi. Optimal exploration-exploitation in a multi-armed-bandit problem with non-stationary rewards. In *NIPS*, 2014.

Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert E Schapire. Contextual bandit algorithms with supervised learning guarantees. In *AISTATS*, 2011.

Avrim Blum and Yishay Monsour. Learning, regret minimization, and equilibria. 2007.

Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Machine Learning*, 5(1):1–122, 2012.

N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.

N. Cesa-Bianchi, Y. Freund, D. P. Helmbold, D. Haussler, R. Schapire, and M. Warmuth. How to use expert advice. *Journal of the ACM*, 44(2):427–485, 1997.

Benjamin Edelman, Sonia Jaffe, and Scott Duke Kominers. To groupon or not to groupon: The profitability of deep discounts. *Marketing Letters*, pages 1–15, 2011.

Yoav Freund and Robert E Schapire. A desicion-theoretic generalization of on-line learning and an application to boosting. In *COLT*, pages 23–37, 1995.

John C Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 148–177, 1979.

Wei-Ning Hsu and Hsuan-Tien Lin. Active learning by learning. In *AAAI*, pages 2659–2665, 2015.

Satyen Kale. Multiarmed bandits with limited expert advice. In *COLT*, pages 107–122, 2014.

J. Langford and T. Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. In *NIPS*, 2007.

Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *WWW*, pages 661–670, 2010.

Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Info and Computation*, 70(2):212–261, 1994.

Odalric-Ambrym Maillard and Rémi Munos. Adaptive bandits: Towards the best history-dependent strategy. In *AISTATS*, pages 570–578, 2011.

H. B. McMahan and M. J. Streeter. Tighter bounds for multi-armed bandits with expert advice. In *COLT*, 2009.

Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.

Adish Singla, Sebastian Tschiatschek, and Andreas Krause. Actively learning hemimetrics with applications to eliciting user preferences. In *ICML*, 2016.

Aleksandrs Slivkins and Eli Upfal. Adapting to a changing environment: the brownian restless bandits. In *COLT*, pages 343–354, 2008.

Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. In *NIPS*, pages 2971–2979, 2015.

P. Whittle. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 1988.

M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, 2003.

Appendix A. Proof of Theorem 1

In this section, we give a proof of the hardness result by discussing a generic and simple setting in which any forecaster suffers a positive average regret.

The setting. Our setting consists of two experts EXP_1 and EXP_2. The first expert EXP_1 has two actions given by $\mathcal{A}_1 = \{a_1, a_2\}$, and the second expert EXP_1 has only one action given by $\mathcal{A}_2 = \{b\}$. The action set of the forecaster, or algorithm ALGO, is given by $\mathcal{A} = \{a_1, a_2, b\}$. The expert EXP_1 plays the HEDGE algorithm (Freund and Schapire, 1995), *i.e.*, the regret rate parameter is $\beta_1 = 0.5$ (see (1)); the expert EXP_2 has only one action to play as in the standard multi-armed bandit with $\beta_2 = 0$ (see (1)). The forecaster knows parameter $\beta = 0.5$ which upper bounds the regret rate of the individual experts.[7]

Loss sequences. Figures 1(*a*), 1(*b*), and 1(*c*) shows the *cumulative* loss sequences L_1, L_2, and L_3 for three different scenarios—the adversary at $t = 0$ uniformly at random picks one of scenarios and uses that loss sequence. These plots show the cumulative losses of the three actions $\mathcal{A} = \{a_1, a_2, b\}$ for three different sequences. For the first scenario with cumulative loss sequences L_1 shown in Figure 1(*a*), we show in Figure 2 the instantaneous losses of the different actions. Figures 2(*a*) and 2(*b*) show the losses of actions for EXP_1; 2(*c*) shows the losses of action for EXP_2.

Model specification. To fully specify the model and Protocol 1, we specify now the feedback vector, and the context over time. The context x^t is constant over time and plays no role in our setting. The experts receive the following feedback when selected: the expert EXP_1 would observe the losses $\{l^t(a_1), l^t(a_2)\}$ when $i^t = 1$; and the expert EXP_2 would observe the loss $\{l^t(b)\}$ when $i^t = 2$.

Execution behavior. We now divide the time horizon T into different slots and discuss the execution behavior of the forecaster and experts over these time slots. Specifically, let us consider the case where the forecaster is facing the sequence L_1 (chosen by adversary with probability $\frac{1}{3}$ at $t = 0$) with cumulative losses shown in Figure 1(*a*) and instantaneous losses of the actions shown in Figure 2. For a clarity of presentation, we shall use $\Delta = 0.01$ as a constant in rest of the proof below.

- $t \in [0, \frac{T}{4}]$: In this time slot, the expert EXP_1 would be selected almost surely by the forecaster and the number of times EXP_2 would be selected is $o(T)$. If this is not the case, then this forecaster would suffer positive average regret on the loss sequence L_3 for third scenario in Figure 1(*c*). [8] The loss incurred by the forecaster at any time t in this time slot is at least 0.5.

- $t \in (\frac{T}{4}, \frac{T}{2}]$: This is the first crucial time slot whereby forecaster's selection strategy would add "blind spots" to the feedback received by EXP_1. The key argument is that in this time slot, the forecaster cannot select expert EXP_1 for more than $\frac{T}{6} + o(T)$ timesteps—if this happens, than this forecaster would have a positive average regret on the loss sequence L_2 for second scenario in Figure 1(*b*). In other words, in this period, the expert EXP_1 has missed seeing the feedback for $\frac{T}{12} - o(T))$ timesteps. Clearly, the loss incurred by the forecaster at any time t in this time slot is at least $0.5 - \frac{3\Delta}{2}$.

7. In fact, this hardness result holds even when considering a powerful forecaster which knows exactly the learning algorithms used by the experts, and is able to see the losses $\{l^t(a_1), l^t(a_2), l^t(b)\}$ at every time $t \in [T]$.

8. Note here that the loss sequence L_3 is exactly equal to L_1 up to time $T/4$. Furthermore, although we have assumed that the setting chosen by the adversary is L_1, we should bear in mind that the forecaster (who can not distinguish between the losses at least up to time $T/4$) should play in a way that it does not suffer positive average regret for L_3.

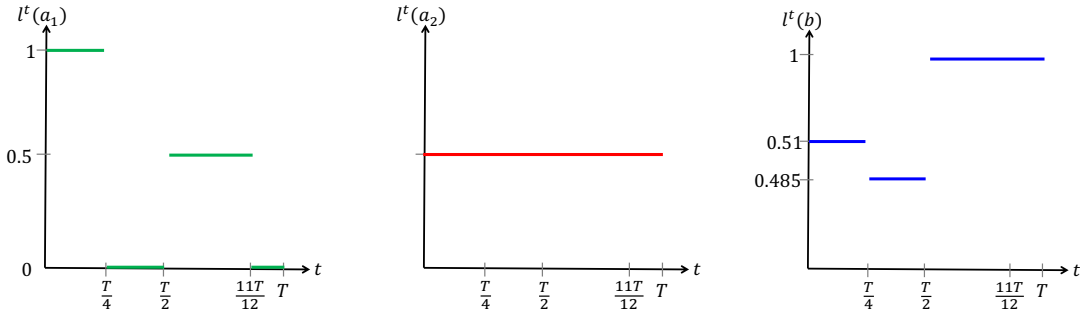

(a) EXP$_1$: Action a_1 (L_1) (b) EXP$_1$: Action a_2 (L_1) (c) EXP$_2$: Action b (L_3)

Figure 2: For the first scenario with the cumulative loss sequences L_1 shown in Figure 1(a), here we show the instantaneous losses of the different actions. Figures 2(a) and 2(b) show the losses of actions for EXP$_1$; 2(c) shows the losses of action for EXP$_2$.

- $t \in (\frac{T}{2}, \frac{11T}{12}]$: In this time slot, the expert EXP$_1$ would be selected almost surely and the number of times EXP$_2$ would be selected is $o(T)$. The loss incurred by the forecaster at any time t in this time slot is at least 0.5.
- $t \in (\frac{11T}{12}, T]$: This is the second crucial time slot which would lead to the positive average regret for the forecaster. We note that the forecaster still does the "right" thing in this time slot, *i.e.*, the expert EXP$_1$ would be selected almost surely and the number of times EXP$_2$ would be selected is $o(T)$. However, the expert EXP$_1$ has missed observing feedback for $(\frac{T}{12} - o(T))$ time steps in the slot $(\frac{T}{4}, \frac{T}{2}]$. Note also that no coordination is permitted between the forecaster and the experts, and hence, EXP$_1$ is not aware of the time steps that it misses the feedback. As a result, at the start of this time slot, the cumulative loss of action a_1 (as perceived by EXP$_1$ based on observed history) is at least $0.5 \cdot \frac{T}{12}$ more than the cumulative loss of action a_2 (as perceived by EXP$_1$ based on observed history). The expert EXP$_1$ who is playing HEDGE algorithm in our setting would select action a_2 almost surely and a_1 would be selected $o(T)$ number of times.

Positive average regret. Let us now compute the regret of the forecaster when experiencing loss sequence L_1 as discussed above. The cumulative loss of the "best expert in hindsight" is given by that of EXP$_1$ always playing action a_1. Based on Figure 2(a), this is given by:

$$\sum_{t\in[T]} l^t(a_1) = 1 \cdot \frac{T}{4} + 0.5 \cdot \big(\frac{11T}{12} - \frac{T}{2}\big) = T \cdot \big(\frac{1}{2} - \frac{1}{24}\big)$$

The cumulative loss of the forecaster as per the execution behavior discussed above can be lower bounded as follows:

$$\begin{aligned}\sum_{t\in[T]} l^t(a^t_{i^t}) &\geq 0.5 \cdot \frac{T}{4} + \big(0.5 - \frac{3\Delta}{2}\big) \cdot \frac{T}{4} + 0.5 \cdot \big(\frac{11T}{12} - \frac{T}{2}\big) + 0.5 \cdot \big(\frac{T}{12} - o(T)\big) \\ &= T \cdot \big(\frac{1}{2} - \frac{3\Delta}{8} - \frac{o(T)}{2 \cdot T}\big)\end{aligned}$$

Hence, the total regret of the forecaster is lower bounded by:

$$\textsc{Reg}(\textsc{Algo}, T) \geq T \cdot (\frac{1}{2} - \frac{3\Delta}{8} - \frac{o(T)}{2 \cdot T}) - T \cdot (\frac{1}{2} - \frac{1}{24}) = T \cdot (\frac{1}{24} - \frac{3\Delta}{8} - \frac{o(T)}{2 \cdot T})$$

Recall that the constant $\Delta = 0.01$, hence the average regret of the forecaster is lower bounded by $\lim_{T \mapsto \infty} \frac{\textsc{Reg}(\textsc{Algo}, T)}{T} \geq \frac{91}{2400}$.

As this sequence L_1 is selected by the adversary uniformly at random with probability $\frac{1}{3}$, this means that the forecaster would suffer a positive average regret. As we discussed above, any forecaster which doesn't have the above-mentioned execution behavior in the timeslot $t \in [0, \frac{T}{4}]$ or $t \in (\frac{T}{4}, \frac{T}{2}]$ would suffer a positive average regret for L_3 and L_2 loss sequences.

Appendix B. Proof of Proposition 2

OCP algorithms. Assume that and expert $\textsc{Exp}_j$ is performing Online Convex Programming (OCP) via greedy projections. We will show that such an algorithm has smooth learning dynamics. Note that OCP has regret of size $O(\sqrt{T})$ (i.e. $\beta_j = 1/2$). Consider the α-OCP algorithm that proceeds according to Algorithm 3. Proving smooth learning dynamics for OCP is equivalent to showing that α-OCP suffers a regret of size $O(\sqrt{T/\alpha})$. More precisely, we have the following Lemma.

Algorithm 3: α-OCP

Problem setting: Convex set S; sequence of convex loss functions $f^t : S \to \mathbb{R}_+$
Parameters: Learning rates η^t for $t \in [T]$
Initialize: $w_0 \in S$ arbitrarily
foreach $t = 1, 2, \ldots, T$ **do**
  $w^{t+1/2} = w^t - \eta^t B^t z^t$ where: (i) $z^t \in \partial f^t(w^t)$, and (ii) random variables B^t are independent Bernoulli with parameter α (i.e. $\Pr(B^t = 1) = 1 - \Pr(B^t = 0) = \alpha$), and (iii) $\eta_t = 1/\sqrt{1 + \sum_{\tau=1}^{t} B^\tau}$
  $w^{t+1} = \text{Proj}_S(w^{t+1/2})$
end

Lemma 4 *Let $\|S\|$ denote the diameter of the convex set S and L denotes an upper bound on the magnitude of the gradient at any time $t \in T$. Then, the expected regret of the α-OCP algorithm is given by*

$$\mathbb{E}[\sum_{t=1}^{T} (f^t(w^t) - f^t(u))] \leq \frac{\|S\|^2}{2} \cdot \sqrt{\frac{T}{\alpha}} + L^2 \cdot \sqrt{\frac{T}{\alpha}} \tag{4}$$

where the expectation is w.r.t. the sequence of Bernoulli random variables B^t for $t \in [T]$.

Proof We can equivalently write the updates in the α-OCP procedure as follows:

$$\begin{aligned} w^{t+1/2} &= w^t - \eta^t B^t z^t \\ &= w^t - (\eta^t \cdot \alpha) \cdot (\frac{B^t}{\alpha}) z^t \end{aligned}$$

$$= w^t - \tilde{\eta}^t \cdot \tilde{z^t}$$

where $\tilde{\eta}^t = (\eta^t \cdot \alpha)$ and $\tilde{z^t} = (\frac{B^t}{\alpha})z^t$. Note that $\mathbb{E}[\tilde{z^t}|w_{1:t}] = \mathbb{E}[z^t] \in \partial f^t(w^t)$. We have:

$$\mathbb{E}\Big[\sum_{t=1}^{T}(f^t(w^t) - f^t(u))\Big] \tag{5}$$

$$= \mathbb{E}\Big[\sum_{t=1}^{T}\mathbb{E}\Big[(f^t(w^t) - f^t(u))|w_{1:t}\Big]\Big] \tag{6}$$

$$\leq \mathbb{E}\Big[\sum_{t=1}^{T}\mathbb{E}\Big[< \partial f^t(w^t), w^t - u > |w_{1:t}\Big]\Big] \tag{7}$$

$$= \mathbb{E}\Big[\sum_{t=1}^{T}\mathbb{E}\Big[< \tilde{z}^t, w^t - u > |w_{1:t}\Big]\Big] \tag{8}$$

$$= \mathbb{E}\Big[\sum_{t=1}^{T}\frac{1}{2\alpha\eta^t}\mathbb{E}\Big[\left\|w^t - w\right\|^2 - \left\|w^{t+1/2} - w\right\|^2 + \eta_t^2\alpha^2\left\|\tilde{z^t}\right\|^2|w_{1:t}\Big]\Big] \tag{9}$$

$$\leq \mathbb{E}\Big[\sum_{t=1}^{T}\frac{1}{2\alpha\eta^t}\mathbb{E}\Big[\left\|w^t - w\right\|^2 - \left\|w^{t+1} - w\right\|^2 + \eta_t^2\alpha^2\left\|\tilde{z^t}\right\|^2|w_{1:t}\Big]\Big] \tag{10}$$

$$= \mathbb{E}\Big[\sum_{t=1}^{T}\frac{\left\|w^t - w\right\|^2}{2\alpha\eta^t} - \frac{\left\|w^{t+1} - w\right\|^2}{2\alpha\eta^t}\Big] + \frac{\alpha}{2}\mathbb{E}\Big[\sum_{t=1}^{T}\eta_t||\theta_t||^2\Big] \tag{11}$$

$$\leq \mathbb{E}\Big[\frac{\left\|w^1 - w\right\|^2}{2\alpha\eta^1} - \frac{||w^{T+1} - w||^2}{2\alpha\eta_T} + \frac{1}{2}\sum_{t=2}^{T}\left\|w^t - w\right\|^2(\frac{1}{\alpha\eta^t} - \frac{1}{\alpha\eta^{t-1}})\Big] + \frac{\alpha}{2}\mathbb{E}\Big[\sum_{t=1}^{T}\eta_t||\tilde{z}^t||^2\Big] \tag{12}$$

$$\leq ||S||^2\mathbb{E}\Big[\frac{1}{2\alpha\eta_T}\Big] + \frac{\alpha}{2}\mathbb{E}\Big[\sum_{t=1}^{T}\eta_t||\tilde{z}^t||^2\Big] \tag{13}$$

Now note that $\mathbb{E}[\eta_t||\tilde{z}^t||^2] = \mathbb{E}\big[\mathbb{E}[\eta_t||\tilde{z}^t||^2|w^{1:t}]\big] = \mathbb{E}\big[\eta_t||z^t||^2/\alpha] \leq L^2\mathbb{E}[\eta_t]/\alpha$. We thus obtain

$$\mathbb{E}\Big[\sum_{t=1}^{T}(f^t(w^t) - f^t(u))\Big] \leq ||S||^2\mathbb{E}\Big[\frac{1}{2\alpha\eta_T}\Big] + \frac{L^2}{2}\mathbb{E}\Big[\sum_{t=1}^{T}\eta_t\Big]. \tag{14}$$

We next recall that $\eta_t = \frac{1}{\sqrt{1+\sum_{\tau=1}^{t}B^\tau}}$. By using the multiplicative Chernoff bound (as B^τ's are Bernoulli random variables) we obtain

$$\Pr(\eta_t \geq \sqrt{2/(t\alpha)}) = \Pr(\sum_{\tau=1}^{t}B^\tau \leq t\alpha/2) \leq \exp(-\frac{t\alpha}{12}).$$

Hence, we obtain $\mathbb{E}[\eta_t] \leq \sqrt{2/(\alpha t)} + \exp(-t\alpha/12)$. Also, due to concavity of the function $h(x) = \sqrt{x}$, we have that $\mathbb{E}[1/\eta_T] \leq \sqrt{T\alpha}$. We finally obtain

$$\mathbb{E}\Big[\sum_{t=1}^{T}(f^t(w^t) - f^t(u))\Big] \leq ||S||^2\sqrt{T/\alpha} + L^2\sqrt{T/\alpha} + \sum_{t=1}^{T}\exp\left(-t\alpha/12\right)$$

$$\leq ||S||^2\sqrt{T/\alpha} + L^2\sqrt{T/\alpha} + 1/(1-\exp(-\alpha/12))$$
$$\leq ||S||^2\sqrt{T/\alpha} + L^2\sqrt{T/\alpha} + 24/\alpha,$$

where the last line is because $1/(1-\exp(-\alpha/12)) \leq 24/\alpha$ for $\alpha \leq 1$. ■

OMD Algorithms. We now consider the case that expert EXP_j is performing an algorithm inside the Online Mirror Descent (OMD) family of algorithms. We assume that the algorithm has a regret of order $O(\sqrt{T})$ for any time horizon T (i.e. $\beta_j = 1/2$). We also assume that the algorithm uses the doubling trick. Consider the standard online learning scenario where at any time $t \in [T]$ a convex function $f^t : S \to \mathbb{R}$ is assigned (S is assumed to be a convex region). The proofs proceeds in 3 steps.

Step 1. α-OMD with a fixed time horizon
We first analyze the algorithm α-OMD given in 4 which is run for a fixed (deterministic) number of steps.

Algorithm 4: α-OMD

Problem setting: Convex set S; sequence of convex loss functions $f^t : S \to \mathbb{R}_+$
Parameters: a link function $g : \mathbb{R}^d \to S$; time horizon T
Initialize: time $\tau = 1$, auxiliary variable $\theta^\tau = \mathbf{0} \in \mathbb{R}^d$
foreach $\tau = 1, 2, \ldots, T$ **do**
  Predict vector $w^\tau = g(\theta^\tau)$
  Update $\theta^{\tau+1} = \theta^\tau - B^\tau z^\tau$ where: (a) $z^\tau \in \partial f^\tau(w^\tau)$, (b) B^τ is an independent Bernoulli random variable with parameter α (i.e. $\Pr(B^\tau = 1) = 1 - \Pr(B^\tau = 0) = \alpha$).
end

Lemma 5 *Let R be a $1/\eta$- strongly convex function over S with respect to a norm $||\cdot||$. Assume that α-OMD is run on the sequence with a link function*

$$g(\theta) = \arg\max_{w \in S}(\langle w, \theta\rangle - R(w))$$

Furthermore, assume that f^t is L-Lipshitz with respect to norm $||\cdot||$. Then

$$\mathbb{E}[\sum_{t=1}^{T}(f^t(w^t) - f^t(u))] \leq R(u)/\alpha + \eta T L^2. \tag{15}$$

Proof For the sake of analysis, we introduce the following slightly modified procedure:

1. Initialize $\tilde{\theta}^1 = \theta^1/\alpha$.
2. At time $\tau = 1, 2, \cdots, T$, let $\tilde{w}^\tau = \tilde{g}(\tilde{\theta}^\tau)$, and $\tilde{\theta}^{\tau+1} = \tilde{\theta}^\tau - \tilde{z}^\tau$. Here, we have $\tilde{z}^\tau = \frac{B^\tau}{\alpha} z^\tau$, and the function $\tilde{g}$ is defined as $\tilde{g}(\theta) = \arg\max_{w \in S}(\langle w, \theta\rangle - R(w)/\alpha)$.

It is straight forward to justify for any $\tau \in [T]$ that $\tilde{\theta}^\tau = \theta^\tau/\alpha$ and $\tilde{w}^\tau = w^\tau$. Also note that $\mathbb{E}[\tilde{z}^\tau | \tilde{z}^{1:\tau-1}] = z^\tau \in \partial f^\tau(w^\tau)$. Hence, the modified procedure $(\tilde{\theta}^\tau, \tilde{w}^\tau)$ is precisely a stochastic

OMD procedure with with link function $\tilde{g}$. By using Theorem 4.1 in Shalev-Shwartz (2011), $\tilde{w}^\tau = w^\tau$, and the fact that $R(\cdot)/\alpha$ is a $1/(\eta\alpha)$-strongly convex function, we obtain

$$\mathbb{E}[\sum_{t=1}^{T}(f^t(w^t) - f^t(u))] \leq \sup_{u\in S} R(u)/\alpha + \eta\alpha \sum_{\tau=1}^{T} \mathbb{E}[||\tilde{z}^\tau||^2].$$

We finally note that

$$\mathbb{E}[||\tilde{z}^\tau||^2] = \mathbb{E}\big[\mathbb{E}[||\tilde{z}^\tau||^2 \,|\, \tilde{z}^{1:\tau-1}]\big] \leq L^2/\alpha.$$

The result of the Lemma is now immediate. ■

Step 2. α-OMD with a random time horizon

From Lemma 5, for $\eta = O(1/\sqrt{T})$, the algorithm α-OMD suffers a $O(\sqrt{T/\alpha})$ regret after any fixed time T. Recall now that at any time the algorithm is only given feedback with independent probability α. We are assuming that the algorithm used by the expert $\textsc{Exp}_j$ is performing the doubling trick, i.e., it runs in blocks whose size get doubled consecutively and within each block the learning rate is fixed. As a result, after the algorithm receives sufficient feedback to finish a block, it restarts OMD and changes the learning rate for the next block (which has twice the size). In order to analayze the regret suffered in each block, we need to consider a slightly different version of α-OMD which stops after a randomly chosen time.

Lemma 6 (α-OMD with a random time horizon) *Assume that we run the α-OMD procedure until the time, call it T_{stop}, such that following stopping criterion has been fulfilled:*

$$\sum_{\tau=1}^{T_{\text{stop}}} B^\tau = M. \tag{16}$$

We the have

$$\mathbb{E}[\sum_{t=1}^{T_{\text{stop}}}(f^t(w^t) - f^t(u))] \leq R(u)/\alpha + \eta M L^2/\alpha + 14L||S||\sqrt{M/\alpha^2}, \tag{17}$$

where $\|S\|$ denote the diameter of the convex set S and L denotes an upper bound on the Lipshitz parameter of all the functions f_t.

Proof We can write

$$\begin{aligned}
&\mathbb{E}[\sum_{t=1}^{T_{\text{stop}}}(f^t(w^t) - f^t(u))] \\
&= \mathbb{E}\big[\sum_{t=1}^{M/\alpha}(f^t(w^t) - f^t(u))\big] - \Big(\mathbb{E}[\sum_{t=1}^{M/\alpha}(f^t(w^t) - f^t(u))] - \mathbb{E}[\sum_{t=1}^{T_{\text{stop}}}(f^t(w^t) - f^t(u))]\Big) \\
&\leq \mathbb{E}[\sum_{t=1}^{M/\alpha}(f^t(w^t) - f^t(u))] + L||S|| \times \mathbb{E}[|T_{\text{stop}} - M/\alpha|],
\end{aligned}$$

where the last step follows from the fact that for any two $u, v \in S$ we have $|f(u) - f(v)| \leq L||S||$. The first term above can be bounded using Lemma 5. We thus need to upper-bound the expected value of $||T_{\text{stop}} - M/\alpha||$. As B^τ's are Bernoulli(α) random variables, we expect that T_{stop} concentrates around M/α. By using the multiplicative Chernoff bound we have

$$\begin{aligned}
\Pr(T_{\text{stop}} \geq M/\alpha + \beta) &= \Pr(\sum_{\tau=1}^{M/\alpha+\beta} B^\tau \leq M) \\
&\leq \Pr(\sum_{\tau=1}^{M/\alpha+\beta} B^\tau \leq (M+\alpha\beta)(1 - \frac{\alpha\beta}{M+\alpha\beta})) \\
&\leq \exp(-\frac{(\alpha\beta)^2}{3(M+\alpha\beta)}).
\end{aligned}$$

Similarly, we can show that

$$\Pr(T_{\text{stop}} \leq M/\alpha - \beta) \leq \exp(-\frac{(\alpha\beta)^2}{3(M-\alpha\beta)}).$$

We thus obtain,

$$\begin{aligned}
\mathbb{E}[|T_{\text{stop}} - M/\alpha|] &\leq \sum_{j=0}^{M/\alpha} \Pr(T_{\text{stop}} \leq M/\alpha - j) + \sum_{j=0}^{\infty} \Pr(T_{\text{stop}} \geq M/\alpha + j) \\
&\leq \sum_{j=0}^{M/\alpha} \exp(-\frac{(\alpha j)^2}{3(M-\alpha j)}) + \sum_{j=0}^{\infty} \exp(-\frac{(\alpha j)^2}{3(M+\alpha j)}) \\
&\leq 2\sum_{j=0}^{\infty} \exp(-\frac{(\alpha j)^2}{3(M+\alpha j)}) \\
&\leq 2\sum_{k=0}^{\infty} \sqrt{M/\alpha^2}\exp(-\frac{Mk^2}{3(M+k\sqrt{M})}) \\
&\leq 2\sqrt{M/\alpha^2}\sum_{k=0}^{\infty} \exp(-\frac{k}{6}) \\
&\leq 14\sqrt{M/\alpha^2}.
\end{aligned}$$

■

Step 3. Putting things together

When the algorithm run by EXP_j is using the doubling trick, for each round (with a block of size M), the algorithm needs to be given M feedbacks (from the forecaster) until it switches to the next round (i.e. it doubles the block-length and restarts the algorithm). Therefore, due to the fact that feedback from the forecaster is sent with independent probability α, the total time needed for the algorithm to switch to the next round is as the one given in Lemma 6. As a result, the regret suffered in the current round is upper-bounded $O(\sqrt{M/\alpha^2})$ (from Lemma 6). Note here that the time spent

in each round to give M feedbacks to the algorithm (i.e. T_{stop} in Lemma 6) is roughly M/α. Now, assume that the total time taken by the algorithm is T. The algorithm (which is given feedback with probability α and plays according to the doubling trick) will be given feedback in $T\alpha$ time units (on average). As a result, it is not hard to see that total regret (after summing up over all the rounds played by the algorithm and using Jensen) becomes $\sqrt{T/\alpha}$. Hence, the proposition is proved also for the OMD algorithms with regret $O(\sqrt{T})$.

Appendix C. Proof of Theorem 3

In this section, we provide the proof of Theorem 3 for the no-regret guarantees of our algorithm LEARNEXP. We follow a step by step approach, beginning with the bounds on external regret of LEARNEXP.

Step 1. Bounds on external regret of LEARNEXP

By directly using the bounds of EXP3 algorithm, cf. Theorem 3.1 from Auer et al. (2002), we can state the following bounds on the external regret of our algorithm LEARNEXP against any expert EXP_k where $k \in [N]$. Note that these bounds given by the *external* regret are only w.r.t. to the post hoc sequence of actions performed and losses observed during the execution of the algorithm

$$\sum_{t=1}^{T} \mathbb{E}\Big[l^t\Big(\pi_{i^t}(x^t, \mathcal{H}^t_{i^t,\text{LEARNEXP}})\Big)\Big] - \mathbb{E}\Big[\sum_{t=1}^{T} l^t\Big(\pi_k(x^t, \mathcal{H}^t_{k,\text{LEARNEXP}})\Big)\Big] \leq c \cdot \eta \cdot T + \frac{(\log N) \cdot N}{\eta} \tag{18}$$

where c is a constant given by $c = e - 1$.

Step2. No-regret and smooth learning dynamics of the experts

We note that during the execution of LEARNEXP in Algorithm 2, we have sparse feedbacks whereby the experts receive feedback instances sporadically at rate defined by $\alpha = \frac{\eta}{N}$, *cf.* Section 4. Hence, by definition, we have $\mathcal{H}^t_{j,\text{LEARNEXP}} \equiv \mathcal{H}^t_{j,\alpha}\ \forall j \in [N]$ where $\alpha = \frac{\eta}{N}$, *cf.* Section 4. By definition, the no-regret smooth learning dynamics of the expert k guarantees:

$$\begin{aligned}\mathbb{E}\Big[\sum_{t=1}^{T} l^t\Big(\pi_k(x^t, \mathcal{H}^t_{k,\text{LEARNEXP}})\Big)\Big] - \mathbb{E}\Big[\sum_{t=1}^{T} l^t\Big(\pi_k(x^t, \mathcal{H}^T_{k,1})\Big)\Big] &\leq \mathcal{O}\Big(T \cdot \Big(\alpha \cdot T\Big)^{\beta_k - 1}\Big) \\ &= \mathcal{O}\Big(\frac{T^{\beta_k} \cdot N^{1-\beta_k}}{\eta^{1-\beta_k}}\Big),\end{aligned} \tag{19}$$

where β_k is the parameter defining the rate of growth of regret, *cf.* Section 4.

Step3. Putting it together

Let us rewrite the regret of the algorithm, copying from Equation 2:

$$\text{REG}(T, \text{LEARNEXP}) \coloneqq \sum_{t=1}^{T} \mathbb{E}\Big[l^t\Big(\pi_{i^t}(x^t, \mathcal{H}^t_{i^t,\text{LEARNEXP}})\Big)\Big] - \min_{j \in [N]} \mathbb{E}\Big[\sum_{t=1}^{T} l^t\Big(\pi_j(x^t, \mathcal{H}^T_{j,1})\Big)\Big] \tag{20}$$

Combining Eq.18 and Eq.19 from above, and using the definition of REG from Equation 20 above, we get:

$$\textsc{Reg}(T, \textsc{LearnExp}) \leq \mathcal{O}\Big(\eta \cdot T + \frac{(\log N) \cdot N}{\eta} + \frac{T^{\beta_k} \cdot N^{1-\beta_k}}{\eta^{1-\beta_k}}\Big) \tag{21}$$

Step4. Optimizing η

Next, we will optimize the value of η in terms of T and N. Note that $\textsc{Exp}_k$ above corresponds to any expert. Hence, let us set $k = j^*$ where j^* corresponds to the best expert $\textsc{Exp}_{j^*}$ that we want to compete against. As per assumptions of the theorem, the best expert indeed has no-regret smooth learning dynamics with $\beta_{j^*} \in [0, 1]$. Stating this in terms $k = j^*$, we can write down the regret as follows:

$$\textsc{Reg}(T, \textsc{LearnExp}) \leq \mathcal{O}\Big(\eta \cdot T + \frac{(\log N) \cdot N}{\eta} + \frac{T^{\beta_{j^*}} \cdot N^{1-\beta_{j^*}}}{\eta^{1-\beta_{j^*}}}\Big) \tag{22}$$

However, note that algorithm doesn't know β_{j^*} and hence cannot directly optimize the value of η. As per the theorem statement, the $\textsc{LearnExp}$ is invoked with input $\beta \in [0, 1]$ such that $\beta \geq \beta_{j^*}$.

Step4.1 Optimizing η for known β_{j^*}, *i.e.*, $\beta = \beta_{j^*}$

To begin with, let us first optimize η for case when $\beta = \beta_{j^*}$. In order to find the optimal dependency of η on T, we set $\eta \sim T^{-z}$, and the value z will be found to minimize the external regret. By this choice of η, the following terms stated as the powers of T appear in (22):

$$\{T^{1-z}, T^{z}, T^{z+\beta_{j^*} \cdot (1-z)}\} \tag{23}$$

Solving for optimal value of z to minimize the power of T in the leading term, we get $z = \frac{1-\beta_{j^*}}{2-\beta_{j^*}}$.

Next, we find the optimal dependency of η on N. Note that, when $\beta = 0$, we have optimal dependency of η on N as $(N \cdot \log(N))^{\frac{1}{2}}$. In general, the optimal dependency of η to N can be found by setting $\eta \sim N^z$, which gives us from (22) the following terms stated as the powers of N, where only the leading terms w.r.t. T are kept:

$$\{N^{z}, N^{(1-\beta_{j^*}) \cdot (1-z)}\} \tag{24}$$

Solving for optimal value of z to minimize the power of T in the leading term, we get $z = \frac{1-\beta_{j^*}}{2-\beta_{j^*}}$.

For any $\beta_{j^*} \in [0, 1]$, we can thus write the optimal value of η as:

$$\eta = T^{-\frac{1-\beta_{j^*}}{2-\beta_{j^*}}} \cdot N^{\frac{1-\beta_{j^*}}{2-\beta_{j^*}}} \cdot (\log N)^{\left(\frac{1}{2} \cdot \mathbf{1}_{\{\beta_{j^*}=0\}}\right)} \tag{25}$$

By keeping only the leading term of T, we can write this as follows:

$$\textsc{Reg}(T, \textsc{LearnExp}) \leq \mathcal{O}\big(T^{\frac{1}{2-\beta_{j^*}}} \cdot N^{\frac{1}{2-\beta_{j^*}}} \cdot (\log N)^{\left(\frac{1}{2} \cdot \mathbf{1}_{\{\beta_{j^*}=0\}}\right)}\big) \tag{26}$$

Step4.2 Optimizing η for unknown β_{j^*}, *i.e.*, $\beta \geq \beta_{j^*}$

When β_{j^*} is not known exactly, and β only upper bounds β_{j^*}, we can still optimize η w.r.t. β to get the same η as stated above, replacing β_{j^*} by β_j (note that $1/(2-\beta)$ is increasing in β). By keeping only the leading term of T, we can write the regret as follows:

$$\textsc{Reg}(T, \textsc{LearnExp}) \leq \mathcal{O}\big(T^{\frac{1}{2-\beta}} \cdot N^{\frac{1}{2-\beta}} \cdot (\log N)^{\left(\frac{1}{2} \cdot \mathbf{1}_{\{\beta=0\}}\right)}\big) \tag{27}$$

This gives us the desired bound stated in Theorem 3.